



Delft University of Technology

Advancing Applications for Artificial-Intelligence-Supported Ambient Control in the Built Environment

Bier, Henriette; Hidding, Arwin; Khademi, Seyran; van Engelenburg, Casper; Alavi, Hamed; Zhong, Sailin

DOI

[10.1080/24751448.2024.2322927](https://doi.org/10.1080/24751448.2024.2322927)

Publication date

2024

Document Version

Final published version

Published in

Technology|Architecture + Design

Citation (APA)

Bier, H., Hidding, A., Khademi, S., van Engelenburg, C., Alavi, H., & Zhong, S. (2024). Advancing Applications for Artificial-Intelligence-Supported Ambient Control in the Built Environment. *Technology|Architecture + Design*, 8(1), 155-164. <https://doi.org/10.1080/24751448.2024.2322927>

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.



Advancing Applications for Artificial-Intelligence-Supported Ambient Control in the Built Environment

Henriette Bier, Arwin Hidding, Seyran Khademi, Casper van Engelenbrug, Hamed Alavi & Sailin Zhong

To cite this article: Henriette Bier, Arwin Hidding, Seyran Khademi, Casper van Engelenbrug, Hamed Alavi & Sailin Zhong (2024) Advancing Applications for Artificial-Intelligence-Supported Ambient Control in the Built Environment, Technology | Architecture + Design, 8:1, 155-164, DOI: [10.1080/24751448.2024.2322927](https://doi.org/10.1080/24751448.2024.2322927)

To link to this article: <https://doi.org/10.1080/24751448.2024.2322927>



© 2024 The Author(s). Published with license by Taylor & Francis Group, LLC.



Published online: 24 May 2024.



Submit your article to this journal [↗](#)



Article views: 99



View related articles [↗](#)



View Crossmark data [↗](#)

Henriette Bier
TU Delft

Arwin Hidding
TU Delft

Seyran Khademi
TU Delft

Casper van Engelenbrug
TU Delft

Hamed Alavi
University of Amsterdam

Sailin Zhong
University of Amsterdam

Advancing Applications for Artificial-Intelligence- Supported Ambient Control in the Built Environment



Ambient intelligence (Aml) relying on electronic devices employing information and communication technology (ICT) and artificial intelligence (AI) embedded in the network connecting these devices tends today to be insufficiently used. This deficiency implies that spaces are uncomfortable and considerable energy dissipates due to distribution losses, excessive or unnecessary climate control of little- and unoccupied spaces, etc.

Building operations are responsible for $\pm 27\%$ of annual carbon dioxide (CO₂) emissions, and infrastructure materials and construction are responsible for an additional $\pm 13\%$ annually; both need to be addressed integratively to meet sustainability goals.^{1,2} This paper addresses this in three AI-supported Aml test simulations of applications focusing on illumination and ventilation systems embedded in the built environment.



◁ Opening Image. Components of a stage integrated with responsive lighting to engage speakers during a symposium at TU Delft. (Credit: TU Delft for all figures unless otherwise noted)

△ Figure 1. Lighting via skylight (left) with unequal distribution (middle) and simulation with reflective rotating panels (right).

Keywords: Ambient Control, AI, Computational Design, Robotics, Interaction, Comfort, Health, CO₂ Reduction, Indoor and Outdoor Climate

Introduction and Context

Ambient Intelligence (Aml) relying on electronic devices employing Information and Communication Technology (ICT) embedded in the network connecting these devices was driven in the last decades by the understanding that sensors and actuators integrated into the environment adapt the respective environment to the users' needs (inter al. Zelkha et al. 1998). Recent advancements in AI enable Aml systems to improve its response to individual requirements, as shown in the case studies presented in this paper.

State-of-the-Art

At its core, Aml refers to environments wherein computing devices are seamlessly integrated. Without AI, these environments may still exhibit some intelligence through predefined rules and simple automation but cannot learn, adapt, or make complex decisions (inter al. Gams et al. 2019).

AI involves various techniques for machinic perceiving, synthesizing, and inferring information. In Aml environments AI is instrumental in creating personalized experiences. By analyzing historical data and user behavior, AI algorithms tailor services and interactions to meet individual preferences, providing a more user-centric and adaptive environment. Furthermore, by analyzing historical data the AI-supported predictive capability allows Aml systems to anticipate user preferences, making proactive adjustments to the environment to enhance user satisfaction. AI also enables Aml systems to continuously learn and improve over time. As these systems gather more data and user feedback, they refine their algorithms and become more adept at meeting user needs and expectations (inter al. Gams et al. 2019). Furthermore, AI contributes to optimizing energy consumption by analyzing data from sensors to make real-time

decisions about lighting and other energy-consuming systems, aiming to reduce energy waste and enhance efficiency (inter al. Lee et al. 2022).

1.2 Contribution

From the plethora of AI-supported approaches for Aml systems, the developed applications employed computer vision (CV) for lighting using digital images from cameras as sensory measures and large machine learning (ML) models,³ aka deep learning (DL), to identify and classify ambient conditions and then react to those (inter al. Nixon and Aguado, 2019) as well as Human Activity Recognition (HAR). Furthermore, for ventilation Autoregressive (AR) and Autoregressive Integrated Moving Average (ARIMA) models using CO₂ data to predict and classify air pollutants have been explored. The proposed integration of such systems into the built environment relies on Design-to-Robotic-Production and Operation (D2RP&O) techniques⁴ that link design with to production and operation of building components (inter al. Bier et al. 2018). While these techniques are not new, their integration with AI-supported approaches for Aml is new and very promising with respect to its potential for increasing users' comfort while reducing energy consumption.

Approach, Methodology, and Results


The case studies addressing lighting and ventilation present what has been achieved so far and indicate the challenges ahead.

2.1 Lighting

Various lighting levels are required for multiple types of work and interaction, e.g., reading, writing, working on the computer, performing, etc. (inter al. Aries 2005). The following two simulation scenarios focused on work and interaction activities.

2.1.1 CV-Supported Optimization of Lighting Conditions

The Computer Vision (CV) case study addressed the problem of



△ Figure 2. Luminance was measured at various days and times for various uses, such as individual and collaborative study (left top and bottom) and training results (right top and bottom) compared to before (middle top and bottom).

poor distribution of daylight in the Technical University (TU) Delft library, which has an open floor plan with flexible furniture configuration. The study explored the potential of 140 reflective panels with adjustable rotation angles to improve lighting conditions by increasing illuminance and reducing substantial differences in light levels that force users' eyes to adapt when moving from one condition to another (Figure 1). The rotation angles of the panels were derived from the best possible light redistribution scenarios based on the possible changes in the furniture layout. To achieve this, several steps were implemented:

2.1.1.1 AI for Panel Configuration

Synthetic images of the space simulated in Rhino Grasshopper⁵ were created to determine the base configurations of different furniture layouts. After capturing images throughout the entire year and creating luminance maps, the necessary panel rotation angles were determined with the Galapagos optimizer to find the most suitable angles of the panels based on the lighting conditions and the type of functional use. These are registered as a 'reference' to look up the best configuration for certain lighting conditions (Figure 2). In this context, the CV model is trained to identify similarities between the actual (live) lighting condition and the closest one in the registry. Once the most similar image regarding lighting condition and functional usage is found then the 'reference' is used to retrieve the best rotation angles for the panels.

2.1.1.2 Image Classification

The same setting with synthetic image generation based on simulation data is used to classify the furniture configurations. This procedure involves a base configuration defined by a

set of cameras for capturing images in the area that requires improvement based on the solar radiance analysis. The dataset was generated by defining certain classes of furniture layouts for individual and collaborative study. The cameras record local variations of these classes, where the position of cameras and furniture pieces varies. This training data augmentation provides additional data while making the training more robust and generalizable to unseen data. The classification labels that correspond to these synthetic training images are the classes indicating the types of furniture layouts. To further increase the size of the training set for the image classification models, variations of images are generated for each class of furniture configurations. Eighty percent of the synthetic dataset is used for training purposes and 20% for testing the model's performance accuracy. A Convolutional Neural Network (CNN) has been trained on the image dataset to assign the images with local variations to the predetermined classes.⁶

2.1.1.3 ML Model

With the data collected in the previous step, the ML model 'learns' from the training datasets to find the best fit. After the training on image datasets presenting the base and varied furniture configurations, the ML model is ready to predict outputs on unseen data. The software for the training involved two pretrained scripts, Resnet18 and VGG16. While Resnet18 takes only the new information from the previous training to the next training layer, which reduces the risk of increasing the inaccuracy in each training layer, VGG16 uses a 3 by 3 pixel filter which is much smaller than the usual filter sizes of other training models. This increases the accuracy of image classification significantly. The advantage of using pretrained models is that the model training can be updated with relatively small datasets to perform successful classification with new but relatively small datasets.

With respect to the training accuracy and validation accuracy (Table 1), currently the model has only been trained on



△ Figure 3. Based on HAR data, responsive lighting changes color, intensity, and on-off pattern.

synthetic datasets, meaning that the value of these results will not be indicative of the real-world performance of the model. Testing on representative datasets will give a true measure of the performance. Both models have a perfect validation accuracy of 1. This is potentially problematic and could point to overfitting, indicating that the model has seen the testing data before or that it has overfitted on the testing data during the development of the model. Another possibility is that the size of the training and testing dataset is too small. In this case, the performance metrics do not reliably indicate the model's ability to generalize unseen data well. Training the model with a larger dataset would help to achieve more reliable metrics for the validation loss and accuracy.

Table 1. Resnet 18 and VGG16 resulting training loss, accuracy, and validation loss for both the furniture classification algorithm and the ML algorithm that optimizes the angles of the panels.

Architecture	RESNET18	VGG16
Training Epochs	24	24
Training Time	34m 18s	15m 18s
Training Loss / Accuracy	0.1503 / 0.9529	0.2348 / 0.9193
Validation Loss / Accuracy	0.0005 / 1.000	0.0042 / 1.000
Best Validation Accuracy	1.000	1.000

Potential next steps include the implementation of the actual library camera footage since there is a difference between the synthetic dataset and the camera data collected from the library. This also applies to the differences between synthetic and actual furniture configurations. This real data will contain noise in the form of people and their belongings, varying lighting during the day, furniture being moved around, etc. Testing on such representative data will give an accurate measure of the robustness of the model. The model can be improved and retrained based on the findings.

Another improvement of the system would be to include predictive control. This involves predicting how the reflections will change over time and using a predictive model to counteract the changing lighting conditions. In an ideal situation, the panels would readjust slowly based on the sun's movement and library camera feedback. The current model could be retrained or extended to meet these requirements.


The traditional approaches have several constraints when comparing the data-intensive approach to traditional light sensor-actuator systems. These systems cannot connect the measured lighting values to the panel angle adjustments, considering the changing lighting conditions and variations in the furniture configurations. Even if a more advanced but not AI-supported system, such as a pre-programmed system that contains a database of the angles of the panels based on the time of the year, would be able to update the panel angles without AI automatically, this system fails when the furniture configurations and weather change over time.

However, the system based on the CNN can update the panels' behavior automatically without human intervention. In this case, the AI-supported Aml approach facilitates the improvement of lighting conditions. It can be applied to other Aml problems involving environmental control performed by wirelessly networked components if the control relies on images—for instance, tracking activities and movement flows as in the following application developed for responsive lighting.

2.2 Responsive Lighting

Responsive lighting was explored to engage speakers with the audience during a symposium at TU Delft (Liu Cheng et al. 2017) by changing light intensity, color, and the on-off rhythm of the light emitting diodes (LEDs) integrated into an adaptive stage (Figure 3). In this case, the AI aspects involve (a) Human Activity Recognition (HAR) and (b) corresponding reactions that promote users' spatial experience via continuous regulation of illumination to activities.

Three reactions were explored: (1) start-up, (2) presentation, and (3) break. In the first, the stage slowly pulsates in one color, suggesting 'awakening,' thus instigating interest in the audience. In the second and third, various interactions were envisioned. First, the stage reacts to the speaker's movements, and the color pattern shifts from start-up to presentation mode. Then, the color pattern changes to the next speaker or break mode as



△ Figure 4. Heterogeneous system architecture facilitates changes of light color, intensity, and on-off pattern (left) of in-building components integrated lighting (right).

soon as the allocated time runs up and according to schedule. In the break, the stage invites the audience to interact with the stage, which instantiates color pattern changes correlated to specific movements. In addition to these automated cause-and-effect modes, the illumination system is equipped with a manual override control.

Integrating the system into the built environment relied on a design-to-production approach that linked the computational design with numerically controlled production machines (Figures 4 and 5). This ensures increased efficiency with machines possibly operating continuously, 24/7, with minimal downtime, contributing to CO₂ reduction. While the two illumination applications are very different in their approach and implementation, both involve AI to adjust illumination by actuating reflective panels and controlling the color, intensity, and on-off pattern of LEDs, respectively. The integration with other environmental aspects and D2RP&O processes has been only partially implemented.

3. Ventilation

Poor indoor air quality has significant adverse impacts on well-being, leading to fatigue, lethargy, headaches, cardiac arrhythmia, and difficulties in attention, memory, and cognitive functioning (Apte and Erdmann 2003; Burge 2004; Erdmann et al. 2002; Fisk 2010; Griffiths and Eftekhari 2008; Seppänen et al. 1999). This is particularly concerning when people spend increasingly more time indoors, estimated at around 90% in developed countries. The issue of high concentrations of indoor air pollutants is especially prevalent in shared spaces

such as meeting rooms and classrooms. For example, a recent study conducted in Switzerland found that two-thirds of the learning spaces in 100 schools exceeded the recommended CO₂ threshold, affecting students' learning capacity (Swiss Federal Office of Public Health 2016). Similarly, a study in the UK demonstrated that unfavorable environmental conditions in offices result in an annual productivity loss of 13 billion pounds (Gorvett 2016).

Furthermore, preliminary studies conducted during the COVID-19 pandemic suggest that improving indoor air quality by increasing the supply of fresh air can help control the spread of the virus in enclosed spaces.

Understanding indoor conditions expands, the field of Indoor Environmental Qualities (IEQ) has begun to explore the opportunities that recent advances in sensing techniques and data science can create to prevent situations of poor air quality in shared spaces. This preventive approach is motivated by two factors: (1) the considerable costs, in terms of both time and productivity, associated with recovering from the consequences of poor indoor air quality, and (2) the potential long-term negative impact on overall well-being resulting from the repeated occurrence of mild health issues such as lethargy, headaches, or other symptoms caused by exposure to inadequate air quality, even during brief periods.

3.1 AI Application for Controlling Indoor Air Quality

Data-oriented methods help predict and prevent poor indoor air quality in shared spaces, as shown in the presented case study, aimed at predicting the carbon dioxide level in the meeting rooms of an office building even before the meeting starts. In the next step, building on the findings of that project, consideration will be given to the ability to modify the space in the building through robotic components that can help reach



TAD 8 : 1

△ Figure 5. Design-to-production of the stage linked the computational design (top left) with numerically controlled production of components (bottom left and right).

the final goal—from the prediction of poor indoor air quality to prevention.

3.1.1 Predicting Indoor Air Quality

The collected air quality data from more than 1000 meeting sessions in an office building was used to examine various ML models that can predict indoor air quality. The concentration of CO₂ indoors, primarily from human respiration, directly correlates with the number of people in a room. However, the CO₂ level is also influenced by factors such as room size, ventilation rate, relative humidity, and outdoor air quality (e.g., Fang et al. 1998). Since measuring all these parameters would require extensive instrumentation of the environment and occupants, the objective is to develop a prediction model that can operate independently of their fluctuations.

Specifically, the aim is to create and compare real-time prediction algorithms that can determine whether the CO₂ level in a room will surpass a predefined threshold based solely on past CO₂ measurements within the same office setting (Alavi et al. 2020).

The application of AR and ARIMA models using CO₂ data collected from shared office spaces and meeting rooms was explored to achieve this. Data was obtained through sensing systems developed in collaboration with an industrial partner. These recorded air pollutant concentrations every five seconds. The Long Short-Term Memory (LSTM), a recurrent neural network architecture, was investigated by

formulating the problem as multiple parallel input and multi-step output scenarios.

The percentage of predicted values that fell within a confidence interval of 30 parts per million (ppm) around the actual value was measured to assess the accuracy of the prediction. This confidence interval was determined based on the technical error range of the sensor. The model's overall accuracy was determined by averaging the accuracies of all prediction instances conducted in a single day of data. These predictions were performed on four devices, with 12 predictions made per hour for 10 hours, excluding the last 20 minutes.

A sliding window approach has been utilized to test the AR and ARIMA models. This involved using an observation buffer to construct the model for predicting the CO₂ concentration in the subsequent Delta T minutes. Various combinations of observation buffer sizes, including 10 and 20 minutes, and Delta T values of 5, 10, and 15 minutes were tested.

Across all the conducted tests, the Autoregressive (AR) model consistently outperformed the other methods in terms of both accuracy and training time. Specifically, when using the AR model with a buffer size of 20 minutes, a prediction accuracy of 97.66% for Delta T = 5 minutes and 87.51% for Delta T = 20 minutes was achieved.

In the next step, the possibility of predicting the future evolution of air quality before a meeting or classroom session is explored (Zhong et al. 2021). Rather than predicting the exact CO₂ concentration level, the objective is to forecast how the CO₂ level will change during the upcoming session. This prediction is based on various parameters such as room size, number of participants, outdoor weather conditions, and time of the day.

A hierarchical clustering analysis of data collected from the meeting sessions held in 26 meeting rooms has been implemented to accomplish this. This analysis allowed the identification of seven distinct patterns of CO₂ evolution.

Each pattern is characterized by an initial value and the rate of increase during the first and second halves of the session (Figure 6). The data for this analysis was obtained from CO₂ sensors developed by an industrial partner, installed on meeting room desks, and recorded CO₂ values every 10 seconds over five months. The study involved more than 300 employees who used the meeting rooms, which were naturally ventilated and varied in size, for sessions typically lasting around one hour.

The goal was to find a combination of external parameters to indicate which of the seven patterns would occur in an upcoming session. Linear Discriminant Analysis (LDA) was employed for this purpose, considering parameters such as room size, number of occupants, indoor conditions (temperature, humidity, light, etc.), outdoor conditions (temperature, humidity, luminosity, wind speed, etc.), time of day, and the concentration level of indoor air pollutants before the session.

The results revealed specific indications in the form of a combination of external parameters that can predict which of the seven patterns of CO₂ evolution is most likely to occur in an upcoming session (Zhong et al. 2021).

3.1.2 Robotic Control Using Predicted Patterns of Air Quality Evolution

Several solutions were envisioned bridging from predicting poor indoor air quality to actions that prevent hazardous conditions. In the test case presented in Section 3.1.1, the goal was to engage the meeting room users to take preventive action using alternative forms of interaction design, namely personal displays (e.g., smartwatches (Zhong et al. 2021)), public devices, and ambient embodied by the meeting room window. In these scenarios, the problem is that interruption is required. In the context of meetings, this can be perceived as intrusive, thus creating a countereffect, with predictions and required countermeasures eventually being ignored. A primary advantage of a robotic automated approach is that it is unobtrusive and functions autonomously without human intervention.

The second possible advantage relates to the autonomous window closing once enough fresh air is provided to lose minimal thermal comfort and minimize the energy needed to maintain thermal comfort. This can be done by simply interfacing the prediction algorithm with the actuation system that controls the ventilation.

However, beyond automating the ventilation system with robotic systems that can control the opening and closing of windows, there is an opportunity to adapt the spatial configuration of the spaces to reshape how the concentration of air pollutants increases in the meeting room.

Previous studies show that among the parameters that determine the level of CO₂ in a shared space, the spatial characteristics of the environment, including the furniture's positioning and the room's geometrical form, are notable. This condition opens an opportunity to examine how, by changing the shape and arrangement of the space (through robotic techniques), one

could reshape the evolution of the concentration of air pollutants and thus postpone the time when preventive action is needed.

Future Steps

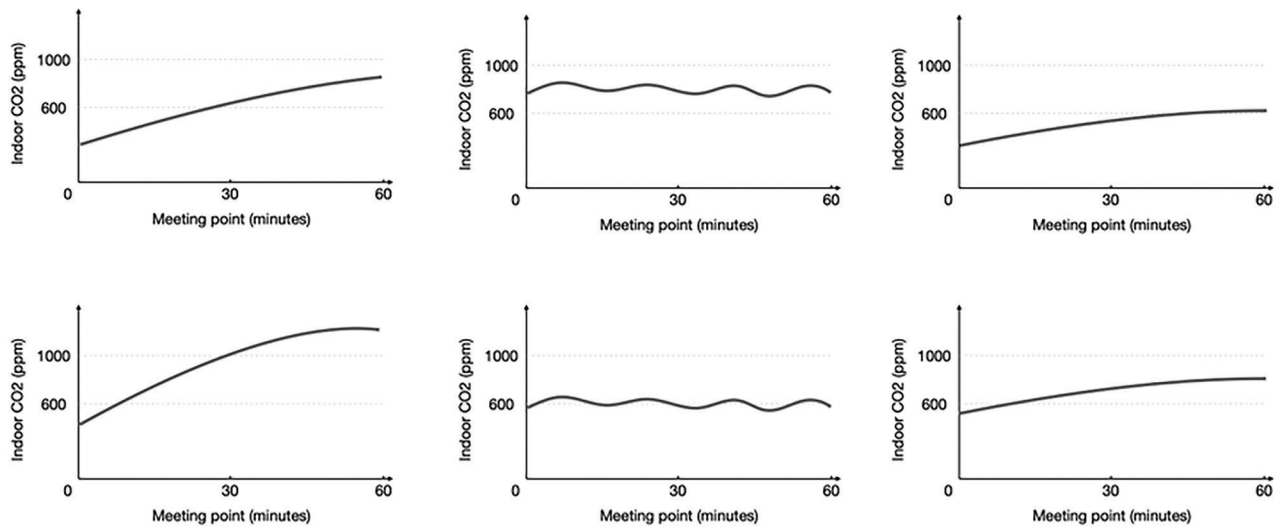
The AI-supported approaches for lighting and ventilation presented will be further advanced and integrated with heating in the future. The overall goal is to determine the indicators that can be applied to tune the environment to the different preferences and needs of occupants for ventilation, lighting, and thermal quality. To adapt to the contextual needs of the users and their activities and achieve an effective embedded Aml, two main challenges have to be addressed at the level of AI development: (a) classifying and detecting the context characterized by parameters such as social and individual activities as well as the subjective perception of comfort assessed by physiological and behavioral signals, and (b) training reinforcements learning models that control the actuators and perpetually correct themselves based on the changes in the contextual parameters. Furthermore, the integration into building components presents challenges for material systems (e.g., walls, floors, etc.) consisting of various subsystems (e.g., sensor-actuators, wiring, etc.) that need consideration from the very start of the D2RP&O process.

In this context, the D2RP&O is part of a larger Design-to-Robotic-Production-Assembly and Operation (D2RPA&O) process that integrates all aspects of building design, construction, and operation from the very beginning of the process and with AI supporting at various stages of the process (inter al. Bier et al. 2022).

The main challenge is to conceptualize Aml as a distributed AI-supported cyber-physical system and embed it into building components.¹ This involves two models:

(a) The human model, which aims to determine approaches that best can be applied to tune the environment to the different preferences and needs of occupants with respect to light, air, and thermal quality. Unobtrusive sensing methods are embedded in the built environment to collect anonymized data about everyone's physiological responses to environmental qualities. For example, visual sensors can log autonomic reactions to lighting conditions such as blinking rate, pupil size change, frowning, and squinting. The inferences from these parameters can be validated against well-established but intrusive methods of predicting human mood, emotion, and comfort, such as on-skin physiological sensing and brain signal loggers. In addition, data about human needs are complemented with information about human conscious preferences and desires through sensing behavioral cues, e.g., interaction with digital and physical building elements.

(b) The built environment model involves the integration of intelligent local control devices into building components, developing reliable indoor environmental indicators and effective D2RP&O mechanisms, i.e., robust control algorithms and fast deploying sensors and actuators, efficient communication protocols for distributed networks, and sustainable embedding procedures. The



△ Figure 6: Schematic illustration of the six distinct ways air quality evolved in the studied meeting room. (Credit: U Amsterdam)

focus is on developing the D2RPA&O process and the 1:1 prototyping of building components with integrated sensor actuators. While the D2RPA&O process implies advancing a reliable design (modeling and simulation) to the production and operation system, the prototyping involves testing and improving building components. By establishing a framework for an integrated approach from 3D to 4D modeling and simulation of indoor-outdoor environments to the D2RPA&O of building components, the knowledge indicating system requirements for dimensions, complexity during installation, and degree of climate control that can be achieved, scalability and life-cycle will be developed.

5. Conclusion

The example simulation cases showcased the use of AI-supported Aml to address lighting and ventilation requirements. Both proved AI's potential to address real-world problems such as Aml with insufficient local control, rendering spaces, if not unhealthy, uncomfortable. Knowing that the integration of AI-supported Aml systems into the built environment using D2RPA&O makes building production and operation more energy efficient, thus reducing CO₂ emissions (inter al. Louis et al. 2014), further advancement is needed as only some of the relevant aspects have been considered so far. The goal is to develop a systematic approach for integrating AI-supported Aml applications into adaptive architecture solutions that can respond to occupants' changing needs and preferences.

Since indoor climate has an impact on outdoor climate and both have a massive impact on humans, advancing AI-supported approaches for ambient control is of great relevance. When acknowledging that building operations are responsible for ±27% of annual CO₂ emissions and infrastructure materials and construction are responsible for an additional ±13% annually, both must be addressed integratively to meet sustainability goals.^{7,8}

While AI-supported lighting and ventilation were explored to a certain degree in the presented case studies, heating still needed to be investigated. Also, they lack integration with D2RPA&O methods, which will be implemented in the next step. The approach to integrate in the built environment AI-supported lighting, ventilation, and heating systems that would automatically adjust to the required environmental conditions based on actively tracked and collected data on outdoor and indoor conditions would ensure control accessible from anywhere via mobile apps while contributing to indoor—and indirectly outdoor—climate improvement.

Acknowledgments

The research profited from the contribution of students and researchers, in particular A. Dosheva, E. Erolsun, G. Park, and S. Knettig, who provided textual and figural representations for the TU Delft library case study.

Data Availability Statement

Data and materials supporting the results presented in their paper are available upon reasonable request.



© 2024 The Author(s). Published with license by Taylor & Francis Group, LLC. This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited. The terms on which this article has been published allow the posting of the Accepted Manuscript in a repository by the author(s) or with their consent.

Notes

1. UN report: <https://www.unep.org/news-and-stories/press-release/co2-emissions-buildings-and-construction-hit-new-high-leaving-sector>.
2. UN-agenda: <https://sdgs.un.org/goals>.
3. ML algorithms build models based on sample data, i.e., training data, to make predictions or decisions.
4. D2RP&O is part of a larger Design-to-Production-Assembly and Operation (D2RPA&O) process that integrates all aspect of building design, construction, and operation from the very beginning of the process.
5. Grasshopper is a visual programming environment that runs within the Rhinoceros 3D application.
6. CNNs are a variant of artificial neural networks that is inspired by biological neural networks able to 'learn' to perform tasks by learning from examples.
7. Cyber-physical systems are composed of physical (i.e., hardware) and cybernetic (i.e., software) systems.
8. UN report: <https://www.unep.org/news-and-stories/press-release/co2-emissions-buildings-and-construction-hit-new-high-leaving-sector>
9. UN-agenda: <https://sdgs.un.org/goals>

References

- Alavi, H., S. Zhong, and D. Lalanne. 2020. "Predictive Models of Indoor Carbon Dioxide Concentration to Prevent Daily Decay of Productivity and Well-Being in Shared Offices." *SmartPhil@ IUI*: 59–68.
- Apte, M. and C. Erdmann. 2002. "Indoor Carbon Dioxide Concentrations, VOCs, Environmental Sensitivity Associations with Mucous Membrane and Lower Respiratory Sick Building Syndrome Symptoms in the BASE Study: Analyses of the 100 Building Dataset." Lawrence Berkeley National Laboratory. <http://escholarship.org/uc/item/3vj5m6wj>.
- Aries, M. 2005. "Human Lighting Demands: Healthy Lighting in an Office Environment." PhD thesis, Technische Universiteit Eindhoven. <https://doi.org/10.6100/IR594257>.
- Bier, H., ed. 2018. *Robotic Building Springer Series in Adaptive Environments*. Springer International Publishing.
- Bier, H., S. Khademi, C. Engelenburg, J. Prendergast, and L. Peternel. 2022. "Computer Vision and Human-Robot Collaboration Supported Design-to-Robotic-Assembly." *Construction Robotics* 6: 251–57. <https://doi.org/10.1007/s41693-022-00084-1>.
- Burge, P. 2004. "Sick Building Syndrome." *Occupational and Environmental Medicine* 61:2: 185–90. <https://doi.org/10.1136/oem.2003.008813>.
- Erdmann, C., K. Steiner, and M. Apte. 2002. "Indoor Carbon Dioxide Concentrations and Sick Building Syndrome Symptoms in the BASE Study Re-Visited: Analyses of the 100 Building Dataset." Technical Report, Lawrence Berkeley National Lab (LBNL), Berkeley, CA.
- Fanger, P. 2000. "Indoor Air Quality in the 21st Century: Search for Excellence." *Indoor Air* 10:2: 63–73.
- Fisk, W. 2010. "CO2 Monitoring for Demand Controlled Ventilation in Commercial Buildings." Lawrence Berkeley National Laboratory. <https://escholarship.org/uc/item/4bk4x479>.
- Gams, M., I. Y-H. Gu, A. Härmä, A. Muñoz, and V. Tam. 2019. "Artificial Intelligence and Ambient Intelligence." *Journal of Ambient Intelligence and Smart Environments* 11: 71–86. <https://doi.org/10.3233/AIS-180508>.
- Govett, Z. 2016. "The Never-Ending Battle over the Best Office Temperature." <https://www.bbc.com/worklife/article/20160617-the-never-ending-battle-over-the-best-office-temperature.com/>.
- Griffiths, M. and M. Eftekhari. 2008. "Control of CO2 in a Naturally Ventilated Classroom." *Energy and Buildings* 40:4: 556–60.
- Lee, D. S., Y. T. Chen, and L. S. Chao. 2022. "Universal Workflow of Artificial Intelligence for Energy Saving." *Energy Reports* 8: 1602–33. <https://doi.org/10.1016/j.egy.2021.12.066>.
- Louis J-N., A. Caló, and E. Pongrácz. 2014. "Smart Houses for Energy Efficiency and Carbon Dioxide Emission Reduction." ENERGY 2014: The Fourth International Conference on Smart Grids, Green Communications, and IT Energy-aware Technologies. SBN: 978-1-61208-332-2.
- Nixon M. and A. Aguado. 2019. *Feature Extraction and Image Processing for Computer Vision*. 4th ed. Academic Press.
- Seppänen, O., W. J. Fisk, and M. J. Mendell. 1999. "Association of Ventilation Rates and CO2 Concentrations with Health and Other Responses in Commercial and Institutional Buildings." *Indoor Air* 9: 226–52. <https://doi.org/10.1111/j.1600-0668.1999.00003.x>
- Swiss Federal Office of Public Health. 2016. Insufficient Air Quality in Two-thirds of the Classrooms in Switzerland. <https://www.admin.ch/gov/fr/accueil/documentation/communiqués.msg-id-74177.html>.
- Zelkha E., B. Epstein, S. Birrell, and C. Dodsworth. 1998. From Devices to 'Ambient Intelligence,' Digital Living Room Conference.
- Zhong, S., H. Alavi, and D. Lalanne, 2020. "Hilo-Wear: Exploring Wearable Interaction with Indoor Air Quality Forecast." In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems*: 1–8.
- Zhong, S., D. Lalanne, and H. Alavi. 2021. "The Complexity of Indoor Air Quality Forecasting and the Simplicity of Interacting with it—A Case Study of 1007 Office Meetings." In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*:1–19.

Henriette Bier is Associate Professor at TU Delft. She leads the Robotic Building Lab, where research focuses on AI-supported robotics integrated into building processes and buildings. Her work is published and exhibited internationally, including in Springer's Adaptive Environments book series, for which she acts as Editor-in-Chief.

Arwin Hidding is a designer and researcher. He graduated cum laude from the TU Delft, and his project was nominated for the Archiprix. He is in academic research and education in the Robotic Building group at TU Delft. The research focuses on 3D prints with programmable properties using different materials or geometries.

Seyran Khademi is an Assistant Professor of Architecture and the Built Environment (ABE) and the Co-Director of AiDAPT lab (AI for Design, Analysis, and Optimization in Architecture and the Built Environment). She works as an Interdisciplinary Researcher between the Computer Vision lab and the Architecture Department at ABE.

Casper van Engelenbrug is a PhD student at TU Delft focusing on understanding visual patterns in floorplan image data. He develops deep contrastive learning frameworks that enable us to learn low-dimensional, task-agnostic representations of architectural drawings. Besides theoretical work, he aims to connect it to the practice by enhancing architectural-specific search engines.

Hamed Alavi's research is focused on the future of human interactive experiences with built environments. Mainly he is interested in the engagement of computer science in the evolution of buildings and urban spaces as they increasingly incorporate artificial intelligence, context-aware automation, and interactivity.

Sailin Zhong is a PhD student at the Human-IST Institute from the University of Fribourg, Switzerland. Zhong's PhD topic is augmenting human perception of comfort in the built environment with interactive AI. Previously, Zhong worked as a Research Assistant at Singapore-ETH center and worked for the Cooling Singapore project for data visualization in Unity.